JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2010, **94,** 241–260 NUMBER 2 (SEPTEMBER)

TOWARD A MECHANICS OF ADAPTIVE BEHAVIOR: EVOLUTIONARY DYNAMICS AND MATCHING THEORY STATICS

J. J McDowell and Andrei Popa

EMORY UNIVERSITY

One theory of behavior dynamics instantiates the idea that behavior evolves in response to selection pressure from the environment in the form of reinforcement. This computational theory implements Darwinian principles of selection, reproduction, and mutation, which operate on a population of potential behaviors by means of a genetic algorithm. The behavior of virtual organisms animated by this theory may be studied in any experimental environment. The evolutionary theory was tested by comparing the steady-state behavior it generated on concurrent schedules to the description of steady state behavior provided by modern matching theory. Ensemble fits of modern matching theory that enforced its constant-*k* requirement and the parametric identities required by its equations, accounted for large proportions of data variance, left random residuals, and yielded parameter estimates with values and properties similar to those obtained in experiments with live organisms. These results indicate that the dynamics of the evolutionary theory and the statics of modern matching theory together constitute a good candidate for a mechanics of adaptive behavior.

Key words: selection by consequences, matching theory, behavior dynamics, complexity theory, computational modeling, concurrent schedules

A quantitative mechanics of adaptive behavior must include an account of behavior that is in equilibrium with conditions in the environment, that is, a statics of behavior, and an account of behavior that is in transition between equilibrium conditions, that is, a dynamics of behavior. Modern matching theory provides an equilibrium account that accurately describes a large and varied body of data from many vertebrate species. An adequate account of behavior dynamics has been more elusive, but a recently proposed computational theory of selection by consequences generates steady state behavior that can be described by at least some of the equations of matching theory (McDowell, 2004; McDowell & Caron, 2007; McDowell, Caron, Kulubekova & Berg, 2008). The purpose of the analyses and experiments reported in the present article was to evaluate this theory more fully, and to assess its viability as a dynamic account that, in combination with modern matching theory, might constitute a reasonably comprehensive mechanics of adaptive behavior. The evolutionary theory was evaluated by examining how well an extensive set of concurrent-schedule data generated by the theory were described by all the applicable equations of modern matching theory. Additional computational experiments were then conducted to extend the domain over which the theory could be evaluated. To motivate this research it is necessary to review modern matching theory and the evolutionary theory of behavior dynamics.

Portions of this research were presented at the Fifth International Conference of the Association for Behavior Analysis in Oslo, Norway, August, 2009. This presentation was supported by a faculty travel grant to the first author from the Institute for Critical International Studies of Emory University. We benefitted greatly from discussions with John Berg, Saule Kulubekova, and especially Marcia Caron.

Address correspondence to J. J. McDowell, Department of Psychology, Emory University, Atlanta, GA 30322 (e-mail: psyjjmd@emory.edu).

doi: 10.1901/jeab.2010.94-241

Behavior at Equilibrium: Modern Matching Theory

Matching theory can be said to exist in two versions. One, classical, version is based on the original matching equation proposed by Herrnstein (1961) and the other, modern, version is based on the power function modification of the original matching equation discussed by Staddon (1968) and Baum and Rachlin (1969), and formalized and studied extensively by Baum (1974, 1979). The two foundational equations are shown in Figure 1 as Equations 1 and 5. In these and all equations in Figure 1 the *B*s represent response rates, the *r*s represent reinforcement rates, numerical subscripts on the rates refer to the two components of a concurrent schedule, and the remaining quantities in the equations

Classic	*Modern*
(1) $\frac{B_1}{B_1+B_2}=\frac{r_1}{r_1+r_2}$	(5) $\frac{B_1}{B_2}=b\left(\frac{r_1}{r_2}\right)^a$
(2) $B=\frac{kr}{r+r_e}$	(6) $B=\frac{kr^a}{r^a+\frac{r_e{}^a}{b}}$
(3) $B_1=\frac{kr_1}{r_1+r_2+r_e}$	(7) $B_1=k\left[\frac{1}{b_{1e}}\left(\frac{r_e}{r_1}\right)^{a_{1e}}+\frac{1}{b_{12}}\left(\frac{r_2}{r_1}\right)^{a_{12}}+1\right]^{-1}$
(4) $B_2=\frac{kr_2}{r_1+r_2+r_e}$	(8) $B_2=k\left[\frac{1}{b_{2e}}\left(\frac{r_e}{r_2}\right)^{a_{2e}}+b_{12}\left(\frac{r_1}{r_2}\right)^{a_{12}}+1\right]^{-1}$
	(9) $\frac{B_1}{B_2}=\frac{1}{b_{2e}}\left(\frac{r_e}{r_2}\right)^{a_{2e}}b_{1e}\left(\frac{r_1}{r_e}\right)^{a_{1e}}$

Fig. 1. The equations of classic (left panels) and modern (right panels) matching theory. The equations are referred to in the text by the numbers that appear in the upper left corner of each panel. The modern equations entail bias parameters and exponents, whereas the classic equations do not. The equations in the first row are the foundational equations of each theory, from which the remaining equations are derived. The modern equations reduce to the classic equations when all bias parameters and exponents equal unity (Equations 5 and 9 both reduce to the ratio form of Equation 1, that is, $B_1/B_2 = r_1/r_2$). Figure and caption are adapted from McDowell's (2005) Figure 1.

are parameters. Herrnstein (1970) showed that Equations 2 through 4 of the classic theory can be derived from the foundational Equation 1 by assuming that a constant summed rate of behavior, k, occurs in all environments. McDowell (1986) showed that Equations 6 through 9 of the modern theory can be derived from the foundational Equation 5 based on the same assumption.

McDowell (2005) argued that all the equations of classic matching theory were false, or at best had limited applicability, including the extensively studied and apparently successful single-alternative hyperbola, Equation 2. His argument was based on a formal analysis of the theory in light of empirical findings, and on evidence indicating that the constant-k assumption of the theory was false. McDowell explained why Equation 2 appeared to describe data well even though it is false, and how the modern theory predicts the violation of the classic theory's constant-k assumption. He also showed that data from a single-schedule experiment that violated the classic theory's constant-k assumption (McDowell & Dallery, 1999) was accurately described by the modern theory, and that data from a concurrent schedule experiment (Dallery, McDowell, & Soto, 2004) was accurately described by all the relevant equations of the modern theory, but not by any of the relevant equations of the classic theory. The latter demonstration entailed fitting the classic theory's Equations 1, 3 and 4 simultaneously, with shared parameters, and fitting forms of the modern theory's Equations 5, 7 and 8 simultaneously, with shared parameters. A subsequent concurrent schedule experiment (Dallery, Soto, & McDowell, 2005) supported McDowell's (2005) conclusion that the classic theory of matching was false, while the modern theory remained tenable as an account of steady-state behavior. In the present article, the modern theory of matching is taken as the correct account of behavior in equilibrium with conditions in the environment. It should be recognized, however, that although the foundational equation of the theory, Equation 5, is almost certainly correct, the remaining equations have not been tested extensively. In particular, the constant-k assumption that is required in order to obtain the absolute response rate forms, Equations 6 through 8, has not been tested rigorously. The relatively unfamiliar equations of modern matching theory, especially Equations 7 and 8, will be explained in more detail later.

Behavior in Transition: Selection by Consequences

McDowell's (2004) selectionist dynamic theory uses Darwinian principles of selection, reproduction, and mutation to cause a population of potential behaviors to evolve under the selection pressure provided by reinforcement from the environment. These principles are implemented by a genetic algorithm that animates virtual organisms created by a computer program. The virtual organisms behave continuously in time, emitting one behavior from the population of potential behaviors each time tick. The emitted behavior can be recorded and studied as if it were the behavior of a live organism. Reinforcement, or selection, in the dynamic theory causes the population of potential behaviors to become more concentrated in the class of behavior targeted for reinforcement, a class analogous to a key peck or a lever press, while variation due to genotypic recombination (McDowell, 2004) and mutation cause behavior to become more distributed among the classes of potential behavior. These opposing forces of variation and selection reach a dynamic equilibrium in a given experimental environment, at which point they generate a roughly constant rate of responding. Additional details of the selectionist theory are provided in the Appendix.

This dynamic theory is not stated in the traditional form of a differential equation, the solution of which produces an equilibrium outcome. Instead, it is stated as a set of low level rules of selection, reproduction, and mutation that operate on a moment-by-moment basis, and that must be applied repeatedly to generate a higher-level time-averaged equilibrium result (McDowell, 2004). This is as an instance of complexity theory, which is a modern alternative to traditional analytic approaches based on the calculus. Some scientists consider complexity theory to be an advance over traditional approaches because it has been able to handle at least some problems that have proved refractory to existing methods (McDowell & Popa, 2009). In complexity theory, higher level outcomes typically cannot be predicted from a theory's low-level rules, and so they are said to be

emergent properties of those rules. To evaluate such a theory it is usually necessary to implement the low-level rules in a computer program, which when run reveals the higher level outcome, and this outcome can then be compared with findings from live experiments.

McDowell and Caron (2007) studied the equilibrium behavior generated by the evolutionary dynamics on single random-interval (RI) schedules. They showed that this behavior was accurately described by the modern theory's single-alternative equation, Equation 6, but not by the classic theory's single-alternative equation, Equation 2. Both equations accounted for large proportions of response rate variance, but the residuals from fits of Equation 2 showed systematic trends, whereas those from fits of Equation 6 did not. McDowell and Caron concluded that, for single RI schedules at least, the selectionist dynamics of the evolutionary theory produced equilibrium behavior that was well described by the modern theory of matching.

More recently, McDowell et al. (2008) studied the behavior of virtual organisms animated by the evolutionary dynamics on concurrent RI RI schedules. At equilibrium the organisms produced behavior that was well described by the power function matching equation, Equation 5, with reasonable and expected values of the parameters, *a* and *b*. They did not, however, conduct the more stringent test of the evolutionary theory recommended by McDowell (2005) for concurrent schedule data. This test engages the entirety of modern matching theory, including its constant-*k* assumption, by fitting forms of Equations 5, 7, and 8 as an ensemble, with shared parameters. Hence, to date, the evolutionary dynamics has been shown to produce behavior at equilibrium that is well described by Equations 5 and 6 of the modern theory of matching, with appropriate and expected parameter values, but that is inconsistent with the corresponding equations of the classic theory (Equations 1 and 2).

In the first part of the present study, the more stringent test of the evolutionary theory recommended by McDowell (2005) was conducted on McDowell et al.'s (2008) extensive set of concurrent schedule data. This was the first test of the evolutionary theory's conformance to forms of the single-alternative Equations 7 and 8 of modern matching theory. Notice that these equations, as well as the corresponding Equations 3 and 4 of the classic theory, are functions of two variables, namely, r_1 and r_2, which are reinforcement rates obtained from the two alternatives of a concurrent schedule. These equations describe surfaces on the two-dimensional reinforcement rate domain. This test of the evolutionary theory, which will be described in detail in the next section, revealed that although the typical concurrent schedule experiment may adequately sample the domain of Equation 5, which consists of reinforcement rate *ratios*, it is likely to sample only a very small portion of the two-dimensional domains of Equations 7 and 8 (and Equations 3 and 4). Hence new experiments were required to test more fully the conformance of behavior generated by the evolutionary dynamics to these equations. The extended domain experiments constituted the second part of the present project.

COMPLETE TEST OF THE EVOLUTIONARY DYNAMICS ON CONCURRENT SCHEDULE DATA

McDowell et al. (2008) studied the behavior of virtual organisms animated by the evolutionary theory on concurrent RI RI schedules. As in every implementation of the theory, parental selection functions (psfs) were used to choose parent behaviors for mating on the basis of their fitness (as described in more detail in the Appendix). The means of these functions can be taken to reflect the impact of each selection event (McDowell, 2004), where smaller means correspond to stronger selection events, and hence are analogous to larger reinforcer magnitudes. The symmetric concurrent schedules in McDowell et al.'s experiment were characterized by psfs having identical means in the two components. Hence, the only difference between the components of the concurrent schedules was the rate of reinforcement delivered by the individual RI schedules.

Before proceeding with the analyses it will be helpful to examine Equations 7 and 8 more carefully. While notationally complicated, these equations are conceptually straightforward. Each expresses the absolute rate of responding in a component of a concurrent schedule as a function of the reinforcement rates obtained from both components. There

are three comparisons in each equation: a comparison of the target component (denoted by the subscript on *B*) with the background (denoted by the subscript, *e*), a comparison of the target component with the other component, and a comparison of the target component with itself. Each comparison consists of a ratio of reinforcement rates, and a bias parameter and exponent that apply to that ratio, as indicated by their subscripts. The three terms inside the square brackets of Equation 7, for example, constitute comparisons of component 1 with the background, component 1 with component 2, and component 1 with itself, the last of which reduces to unity. As noted, each reinforcement rate ratio has an appropriately subscripted bias parameter and exponent. Similarly, the three terms inside the square brackets of Equation 8 constitute comparisons of component 2 with the background, component 2 with component 1, and component 2 with itself. It is important to recognize that these equations were not assembled on any *a priori* logical or theoretical grounds, they are simply the mathematical consequence of applying the constant-*k* assumption to Equation 5, and then working through the algebra (McDowell, 1986). The calculation itself produced the comparisons that appear in Equations 7 and 8.

McDowell (2005) noted that the equations of modern matching theory could be simplified for fitting by setting the three exponents they entail to the same value. He also pointed out that not all the parameters in the equations can be estimated independently. It is possible, however, to assemble composite parameters from the parameters native to the equations that do permit independent estimation. With these composite parameters, and the recommended exponent simplification, McDowell (2005) showed that Equations 5, 7, and 8 become

$$\frac{B_1}{B_2} = \frac{c_{2e}}{c_{1e}}\left(\frac{r_1}{r_2}\right)^a, \qquad (5')$$

$$B_1 = k\left[c_{1e}\left(\frac{1}{r_1}\right)^a + \frac{c_{1e}}{c_{2e}}\left(\frac{r_2}{r_1}\right)^a + 1\right]^{-1}, \qquad (7')$$

and

$$B_2 = k\left[c_{2e}\left(\frac{1}{r_2}\right)^a + \frac{c_{2e}}{c_{1e}}\left(\frac{r_1}{r_2}\right)^a + 1\right]^{-1}, \qquad (8')$$

with composite, independently estimable, parameters,

$$c_{1e} = \frac{r_e^a}{b_{1e}},$$

and

$$c_{2e} = \frac{r_e^a}{b_{2e}}.$$

McDowell also showed that with the recommended exponent simplification,

$$b_{12} = \frac{b_{1e}}{b_{2e}} = \frac{c_{2e}}{c_{1e}},$$

and, importantly, that Equation 9 in Figure 1 reduces to Equation 5′ above.

Equations 5′, 7′ and 8′ constitute a complete statement of modern matching theory. Equation 5′ is the well known power function matching equation, and Equations 7′ and 8′ are single-alternative equations that are required by the theory to describe responding in each component of a concurrent schedule. Note that both Equation 7′ and Equation 8′ reduce to Equation 6 when $r_2 = 0$ in Equation 7′ and $r_1 = 0$ in Equation 8′, which means they are both forms of the canonical absolute response rate equation, Equation 6. Note also that Equations 7′ and 8′ engage the constant-*k* assumption of the modern theory inasmuch as this assumption is required for their derivation. The three equations entail a total of four free parameters, namely, *k*, which appears in both Equations 7′ and 8′, and *a*, c_{1e}, and c_{2e}, which appear in all three equations.

Method

In Phases 1 and 3 of McDowell et al.'s (2008) experiment, four parental selection function means that were identical in the two components (namely, 20, 40, 60, and 80) were studied at each of 10 mutation rates (0.5, 1, 2, 3, 5, 7.5, 10, 12, 20, and 50%). Details of the theory's implementation in these phases are given in the Appendix. The rate of mutation in the evolutionary theory is probably best conceptualized as an organismic variable. For example, McDowell et al. reported that at low mutation rates behavior generated by the theory tended to be perseverative, getting stuck in particular classes, and was relatively insensitive to obtained reinforcement rate ratios. At high mutation rates, on the other hand, behavior tended to be

impulsive, jumping from class to class, and showed a similar relative insensitivity to the reinforcement rate ratios.

McDowell et al.'s (2008) Table 3 summarizes the 40 conditions they arranged. Each entailed 11 concurrent schedules with independent RI components ranging from RI 20 (time ticks) to RI 120. Each concurrent schedule was in effect for 20,500 generations, or time ticks, and each of the 40 conditions was replicated 5 to 20 times. The 40 conditions (psf mean pair × mutation rate), 11 concurrent schedules per condition, and 5 to 20 replications of each condition, constituted about 72 million generations of responding. In their analyses, they fitted Equation 5, the power function matching equation, to the reinforcement and response rate ratios from each replication of a condition, resulting in 330 fits of that equation. For the present analyses, the reinforcement and response rates were first averaged over the 5 to 20 replications of each condition, yielding 40 averaged data sets. All analyses were then conducted on these averages.

Ten fits of Equations 5′, 7′ and 8′ were conducted with these data, one for each mutation rate. If mutation rate is conceptualized as an organismic variable, then this is analogous to fitting the equations to data from 10 subjects. At each mutation rate, the three equations were fitted simultaneously to the data from all four psf mean pairs. Hence each fit entailed 4 of the 40 conditions of the experiment. Assuming that the psf mean corresponds to reinforcer magnitude, the constant-k assumption of matching theory requires k to remain constant across the four psf mean pairs, and hence it is theoretically necessary to conduct this fit with just one shared k value. Extending the ensemble fitting method developed by McDowell (2005) to this more complicated situation, a normalized residual sum of squares (RSS) was calculated for each pair of psf means:

$$NormalizedRSSSum_i = \frac{RSS_{7'}}{SS_{7'}} + \frac{RSS_{8'}}{SS_{8'}} + \frac{RSS_{5'}}{SS_{5'}}, \qquad (10)$$

where SS represents the total sum of squares, the numerical subscripts refer to the three equations, and the index, i, refers to the i^{th} psf mean pair. The four sums generated in this way were then themselves summed to produce the quantity to minimize in the ensemble fit, namely,

$$TotalNormalizedRSSSum = \sum_{i=1}^{4} NormalizedRSSSum_i, \qquad (11)$$

where the summation extends over the four psf mean pairs.

This ensemble fit entailed three sources of variance (B_1, B_2, and B_1/B_2) at each psf mean pair, yielding a total of 12 sources of variance across the four psf mean pairs that were fitted simultaneously. For each fit, the quantity defined by Equation 11 was minimized by adjusting one k parameter, and one a, c_{1e}, and c_{2e} parameter for each psf mean pair. The single k parameter was shared by eight equations in the ensemble fit, and each of the a, c_{1e}, and c_{2e} parameters were shared by the three equations at each psf mean pair. Hence, for every ensemble fit, 12 sources of variance were fitted simultaneously by adjusting 13 parameters to minimize the quantity defined by Equation 11.

To summarize and recapitulate the analysis, consider Figure 2, which is a schematic diagram of a data set to which Equations 5′, 7′, and 8′ may be fitted as an ensemble. Each square represents a source of variance (reinforcement and response rates, or ratios, from 11 concurrent schedules). As discussed earlier, there are three sources of variance at each psf mean pair and each has its own RSS (from the fitted equation) and SS. For the ensemble fit, the 12 individual quotients, RSS/SS, were summed (Equation 11), and this sum was minimized by adjusting one k shared across all psf mean pairs, and one a, c_{1e}, and c_{2e} shared by the three sources of variance at each psf mean pair. Because k was not permitted to vary across any of the individual sources of variance, and because the other parameters were not permitted to vary across the three sources of variance at each psf mean pair, these fits were highly restrictive. They constituted a stringent test of how well the behavior produced by the evolutionary dynamics conformed to the steady-state equations of modern matching theory.

To compare the modern theory with the classic theory, ten analogous ensemble fits of

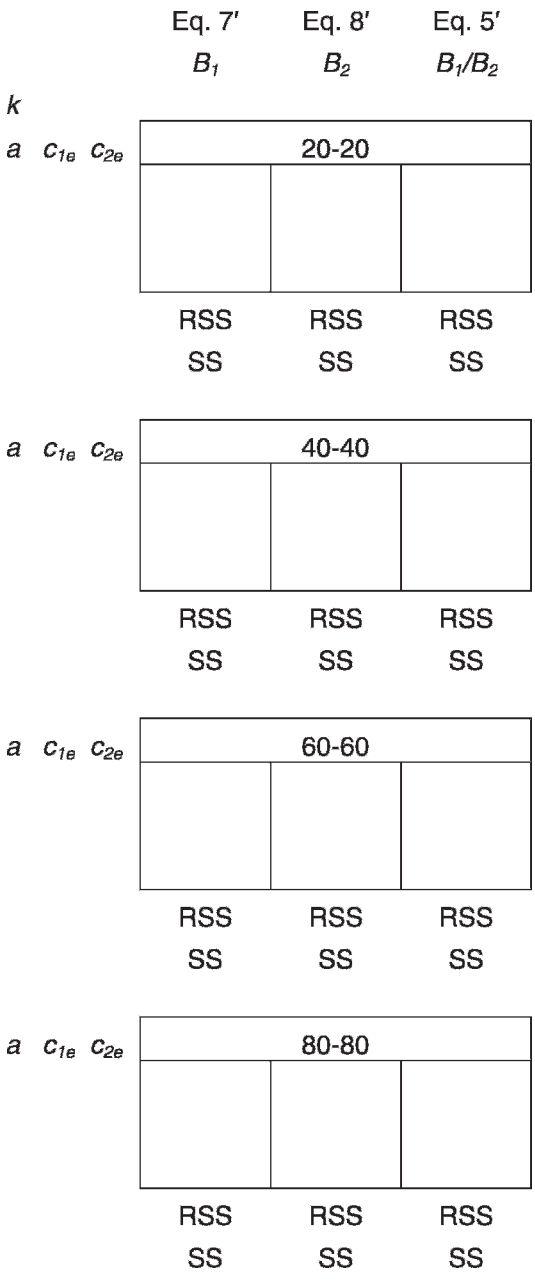

Fig. 2. Schematic diagram of a data set to which the three equations of modern matching theory may be fitted as an ensemble. Twelve sources of variance (reinforcement and response rates from concurrent schedules), which are represented by squares, are fitted simultaneously by adjusting a single k, shared across all psf mean pairs (20-20, 40-40, etc.), and one a, c_{1e}, and c_{2e}, at each psf mean pair, shared by the three equations at that mean pair. The 13 parameters are adjusted simultaneously to minimize the sum of the twelve individual RSS/SS ratios (Equation 11), one ratio calculated from each source of variance. A similar schematic diagram can be drawn to illustrate how the three equations of classic matching theory are fitted as an ensemble to a set of data. The same quantity is minimized, but by adjusting a single k, shared across all psf mean pairs, and one r_e at each psf mean pair, shared by the three equations at that mean pair.

Equations 1, 3, and 4 were carried out on this data set. At each mutation rate, the quantity defined by Equation 11 was minimized by adjusting a single k parameter and four r_e parameters, one r_e for each psf mean pair. Hence for each ensemble fit, 12 sources of variance were fitted simultaneously by adjusting five parameters to minimize the quantity defined by Equation 11. Again, these were highly restrictive fits and constituted a stringent test of the evolutionary dynamics' consistency with the steady-state requirements of classic matching theory.

Results and Discussion

Percentages of variance accounted for (%VAF) were calculated separately for each source of variance. This permitted the quality of the fit to be examined for each source separately. Hence, 12 %VAFs were calculated for each of the 10 ensemble fits of the modern theory equations, and for each of the 10 ensemble fits of the classic theory equations. The median %VAFs for the three equations of the modern theory and the three equations of the classic theory across psf mean pairs and mutation rates are plotted in Figure 3. The error bars span the interquartile range. As this plot shows, the %VAFs for fits of the modern theory were uniformly large, while for the classic theory they were more variable, although the median %VAF was at or above a respectable 80% for the three equations of the classic theory. For both theories the %VAFs tended to be smaller at the lowest and the highest mutation rates, and for the classic theory they were often negative at these rates. Negative %VAFs indicate that the data varied more from the fitted function than from their own mean. Importantly, the %VAFs for the modern theory were greater than the corresponding %VAFs for the classic theory in all individual comparisons of the 120 sources of variance.

Of course it is not surprising that higher %VAFs characterized the ensemble fits of the

←

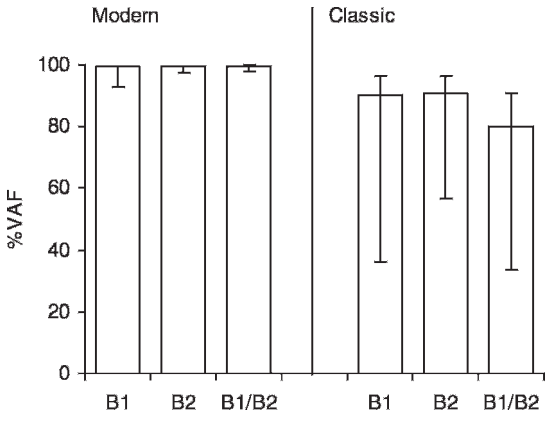

Fig. 3. Median percentages of variance accounted for (%VAF) by Equations 7′, 8′, and 5′ of the modern theory, and Equations 3, 4, and 1 of the classic theory of matching. The *x*-axis labels identify the dependent variable of each equation. The error bars span the interquartile range.

modern theory inasmuch as they entailed 13 free parameters compared to the 5 free parameters for each ensemble fit of the classic theory. The benefit of the higher %VAFs can be weighed against the cost of the additional parameters by comparing the corrected Akaike Information Criterion (AIC_c) calculated for the modern and classic theories. This quantity is calculated from the number of data points and the number of parameters involved in the fit, and the *RSS* of the best fit (Motulsky & Christopoulos, 2004). The AIC_c increases with each of these quantities. When comparing competing theories, the theory with the smaller AIC_c is considered the better theory. Because a more complicated account has more parameters, it will have a larger AIC_c. But if the *RSS* it leaves is sufficiently smaller than the *RSS* for the less complicated account, then its AIC_c may be reduced to a value smaller than the AIC_c for the less complicated account. This would mean that the benefit of the reduced *RSS* more than made up for the larger number of parameters.

One AIC_c was calculated for each of the 10 ensemble fits of the modern theory and each of the 10 ensemble fits of the classic theory. Equations 1 and 5′ were omitted from these calculations because it is well known that Equation 1 generally does not describe concurrent schedule data well, and so including this equation would seem to give an unfair advantage to the modern theory. Each AIC_c therefore entailed 88 data points (one B_1 and one B_2 for each of 11 concurrent schedules × 4 psf mean pairs) and the RSS used for each AIC_c was the sum of the eight relevant individual RSSs (1 for each concurrent schedule component × 4 psf mean pairs). The schematic diagram in Figure 2 may be helpful in understanding how these quantities were calculated.

The absolute values of the AIC_cs for the ensemble fits were large because of the large number of data points and the correspondingly large number of residuals that contributed to the RSSs used in their calculation. The *differences* between the AIC_cs for the modern and classic theories ranged from 197 to 445 for the ten comparisons, in every case favoring the modern theory. These differences were extremely large because the RSSs left by the modern theory were much smaller, by at least an order of magnitude, than those left by the classic theory. These AIC_c differences indicated overwhelmingly that the modern theory provided the better account of the data generated by the evolutionary theory, even given the additional parameters it entailed. Evidently, the much smaller RSSs left by the modern theory more than made up for the larger number of parameters.

Residuals left by the ensemble fits of the modern and classic theory equations were examined by plotting the standardized residuals against the dependent variables predicted by the best fitting equations. For each equation the residuals were pooled across the four psf mean pairs, generating 30 sets of residuals for the modern theory (3 equations × 10 ensemble fits), and 30 sets of residuals for the classic theory. A cubic polynomial was fitted to each set of residuals (cf. McDowell & Caron, 2007). Significant ($\alpha = .05$) cubic polynomial trends were found for 7 of the 30 sets of the modern theory's residuals, and for all 30 sets of the classic theory's residuals. Applying a binomial test, similar to the one recommended by McDowell (2004), the binomial probability of identifying 7 of 30 sets of residuals as nonrandom under the null hypothesis that all were random was greater than .05, indicating that these residuals can be considered random. In contrast, the binomial probability of identifying 30 of 30 sets of residuals as nonrandom under the null hypothesis that all were random is zero to many decimal places, indicating that the residuals for the

classic theory cannot be considered random. Moreover, the effect sizes (R^2) of the trends in the classic theory's residuals were quite large, ranging from 0.49 to 1.00, with a median of 0.88. In contrast, the effect sizes for the modern theory's residuals ranged from 0.01 to 0.35 with a median of 0.10.

The parameters of the best fitting equations of the modern theory and the classic theory are listed in Table 1 for each mutation rate and psf mean pair. The c_{2e}/c_{1e} parameters in the table, which are the bias parameters in Equation 5 (compare Equations 5 and 5′), were calculated from the estimated c_{1e} and c_{2e} parameters. The exponents from the ensemble fits are plotted in the top panel of Figure 4. The legend in the center panel applies to both the top and center panels. As shown in the figure, the exponents increased from very low values at low mutation rates, to a value of about 0.8 at mutation rates around 10%, and then decreased again to very low values as the mutation rate increased further. In addition to this marked change as a function of mutation rate, the exponents also tended to be larger for smaller psf mean pairs (stronger selection events), an effect that was more pronounced at the low and high mutation rates than at the intermediate rates. Exponents greater than or equal to 0.7 are plotted in the center panel of Figure 4. This constitutes a range of exponents, from 0.70 to 0.86, that is more or less typical of live organisms responding in laboratory experiments. Notice that the *x*-axis in this panel extends to a mutation rate of only 20%. This plot shows that stronger selection events (smaller psf mean pairs) produced exponents more or less typical of live organisms over a larger range of mutation rates than did weaker selection events.

The top two panels of Figure 4 can be compared to the top two panels of McDowell et al.'s (2008) Figure 4, which show exponents for the same data, but obtained from traditional fits of Equation 5 only, rather than from the ensemble fits used here. The plots are very similar, with the exception that in the present Figure 4, the exponents fell much more markedly at low mutation rates than in McDowell et al.'s Figure 4. This is probably due to McDowell et al. fitting Equation 5 to each repetition of a condition separately, rather than to the average of the repetitions, as was done here. At very low mutation rates, the perseverative behavior of the virtual organisms often resulted in zero reinforcement and response rates in one or the other component of the concurrent schedules. Schedules with zero rates could not be used in the fits for obvious reasons. But when the repetitions were averaged first, as was done for the present analyses, the zero rates were eliminated because the perseveration often occurred in different components from repetition to repetition. Hence, because the present method of averaging across repetitions makes it unnecessary to discard data, the present Figure 4 evidently represents the workings of the evolutionary dynamics at low mutation rates more accurately than does McDowell et al.'s Figure 4.

The bottom panel of Figure 4 shows how the bias parameter, b_{1e}, in Equations 7 and 8 behaves as a function of the psf mean in a component. Recall that smaller psf means may correspond to larger reinforcer magnitudes. Hence bias in favor of a concurrent schedule component relative to the unchanging background alternative should *decrease* as the psf mean in the component increases. From the earlier discussion of the composite *c* parameters, it follows that

$$\frac{b_{1e}}{r_e} = \frac{1}{\sqrt[a]{c_{1e}}}. \tag{12}$$

Hence, b_{1e}/r_e can be calculated from the estimates of a and c_{1e} that were obtained from the ensemble fit. If the rate of background reinforcement is unchanging, which evidently is required by matching theory (but cf. Soto, McDowell, & Dallery, 2005), then the behavior of b_{1e}/r_e must be due solely to changes in b_{1e}. Hence for the selectionist dynamic theory to be consistent with matching theory, b_{1e}/r_e must decrease as the psf mean in a component increases, that is, as the selection events in that component become weaker. The plots in the bottom panel of Figure 4 show that this outcome was in fact obtained. In other words, the behavior of the c_{1e} parameter, and hence of the b_{1e} parameter, in these fits was consistent with matching theory's interpretation of bias. An identical result can be shown for the quantity b_{2e}/r_e, which can be calculated from the estimates of a and c_{2e} using an expression analogous to Equation 12. The plots in the bottom panel of Figure 4 also show

Table 1

Parameters of the best fitting equations of modern and classic matching theory.

		Modern					Classic	
psf mean pair		k	a	c_{1e}	c_{2e}	c_{2e}/c_{1e}	k	r_e
	0.5%							
20-20		343.80	0.62	0.17	0.18	1.09	164.49	0.00
40-40			0.49	1.85	2.03	1.10		0.00
60-60			0.33	2.54	2.84	1.12		0.19
80-80			0.22	3.49	3.33	0.96		1.88
	1%							
20-20		362.63	0.68	1.06	1.11	1.04	231.28	0.00
40-40			0.60	3.19	3.15	0.99		1.40
60-60			0.44	3.90	4.03	1.03		4.35
80-80			0.46	5.66	5.70	1.01		7.05
	2%							
20-20		364.91	0.74	2.51	2.48	0.99	234.45	0.00
40-40			0.68	5.11	5.03	0.98		2.84
60-60			0.60	6.63	6.70	1.01		6.25
80-80			0.56	8.27	8.26	1.00		9.52
	3%							
20-20		357.17	0.78	3.64	3.70	1.01	220.37	0.00
40-40			0.73	6.92	6.83	0.99		3.46
60-60			0.66	8.74	9.12	1.04		7.31
80-80			0.63	11.46	11.36	0.99		11.20
	5%							
20-20		366.85	0.83	7.30	7.24	0.99	224.47	2.33
40-40			0.77	11.46	11.51	1.00		7.08
60-60			0.75	15.19	15.16	1.00		11.42
80-80			0.74	19.72	19.21	0.97		16.27
	7.5%							
20-20		403.61	0.82	12.95	12.69	0.98	212.83	4.55
40-40			0.80	19.07	19.27	1.01		9.89
60-60			0.78	24.72	24.84	1.01		15.07
80-80			0.72	28.09	28.15	1.00		20.81
	10%							
20-20		408.90	0.86	17.77	18.24	1.03	192.68	5.17
40-40			0.81	25.42	24.97	0.98		11.46
60-60			0.76	30.13	29.54	0.98		16.76
80-80			0.71	33.82	33.34	0.99		22.18
	12%							
20-20		426.65	0.86	22.74	22.65	1.00	169.27	4.96
40-40			0.81	30.83	30.72	1.00		10.96
60-60			0.75	35.18	34.20	0.97		15.70
80-80			0.66	35.18	35.53	1.01		20.91
	20%							
20-20		479.11	0.81	36.01	35.93	1.00	89.89	0.43
40-40			0.70	41.01	40.95	1.00		4.22
60-60			0.57	38.88	39.16	1.01		7.39
80-80			0.45	35.91	35.66	0.99		9.78
	50%							
20-20		414.45	0.52	30.97	30.82	1.00	45.56	0.00
40-40			0.33	26.76	26.71	1.00		0.00
60-60			0.21	23.38	23.40	1.00		1.07
80-80			0.16	22.20	22.25	1.00		1.57

Note. The parameter, c_{2e}/c_{1e}, was calculated from the estimates of c_{1e} and c_{2e}. The spanner heads identify the mutation rate; psf stands for parental selection function

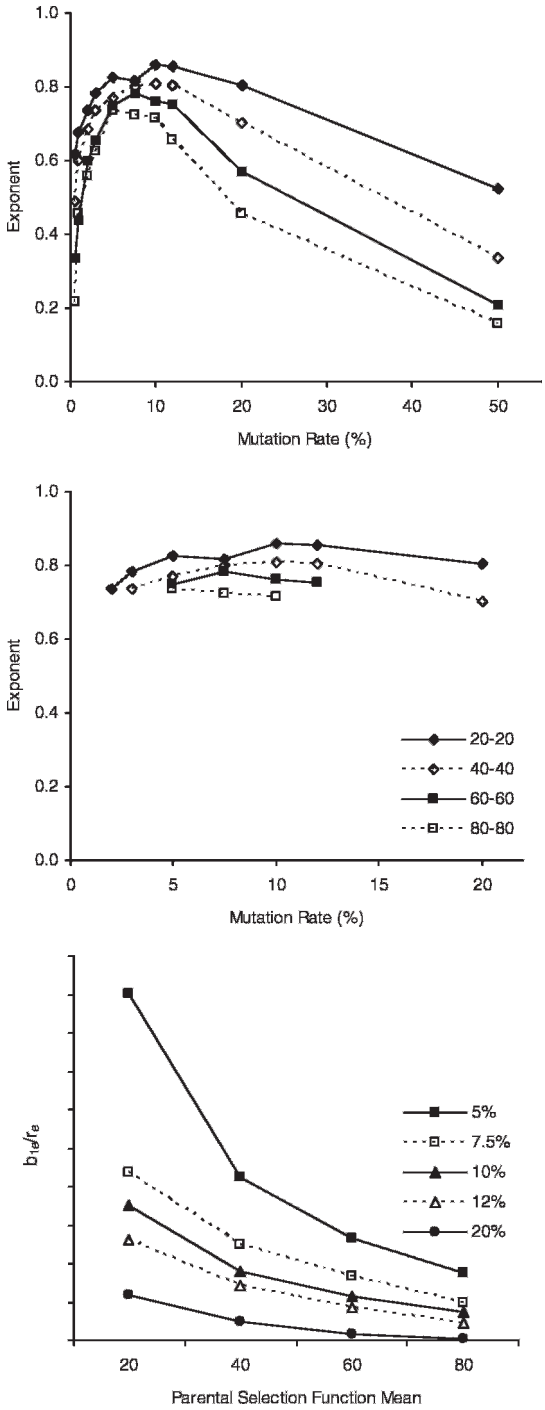

Fig. 4. Top and center panels: exponents from Table 1 plotted as a function of mutation rate. In the center panel, only exponents ≥ 0.70 are plotted. Notice that the *x*-axis in this panel extends to a mutation rate of only 20%. The legend in the center panel applies to both the top and center panels. Bottom panel: plot of a quantity related to bias, viz., b_{1e}/r_e, as a function of the psf mean in a concurrent schedule component. The legend refers to mutation rates. At all mutation rates, including those not plotted in the panel, bias for the concurrent schedule component over the background alternative decreased as the selection event in the component weakened, that is, as the psf mean became smaller.

that bias favoring the concurrent schedule component over the background alternative decreased with increasing mutation rate, and this was the case for both b_{1e} and b_{2e}.

Finally, in these symmetrical concurrent schedules, the bias parameters, c_{2e}/c_{1e}, in Equation 5′ must equal unity according to matching theory. Because the only difference between the components of the concurrent schedules were the rates of reinforcement, there should be no response bias favoring one component over the other. The values listed in Table 1 show that these parameters did in fact vary around unity, with the variability being rather marked at the lowest mutation rate.

Overall, the %VAF and AIC_c comparisons, the residual trends, and the parametric behavior of the fitted equations in these analyses support the conclusion that the evolutionary theory generated steady-state behavior that was consistent with the modern theory of matching, but that was inconsistent with the classic theory. One potential problem with these analyses, however, is that experiments designed to adequately sample the reinforcement *ratio* domain of Equation 5 may not adequately sample the two-dimensional absolute reinforcement rate domains of Equations 3, 4, 7, 8, and related forms (e.g., Equations 7′ and 8′). This problem is illustrated in Figure 5, which shows two-dimensional absolute reinforcement rate floors in the top and bottom panels. In studies of the evolutionary theory time is typically divided into 500-tick blocks, which are taken as standard units, analogous perhaps to hours (McDowell, 2004; McDowell et al., 2008). Because a reinforcer can be delivered during any of the these times ticks, reinforcement rates in principle can vary from 0 to 500 reinforcements per 500-tick block. The concurrent schedules studied here arranged RI schedules with values between 20 and 120 and sampled a wide range of scheduled reinforcement rate ratios, ranging from 0.17 to 6.00. But RI values in the

components of these schedules translate to absolute scheduled rates of reinforcement ranging from only 4.17 to 25 reinforcements per 500-tick block. These scheduled rates are plotted as filled symbols in the top panel of Figure 5. Evidently, they cover only a minuscule portion of the reinforcement rate floor, sampling only about ¼ of 1 percent of the two-dimensional domains of the absolute response rate equations. It may be that findings from so limited a sampling of the reinforcement rate domain are inaccurate or misleading.

It is of course possible to move further out into the domain by decreasing the component RI values, as illustrated in the bottom panel of Figure 5. All reinforcement rate contours in this panel sample the range of scheduled reinforcement rate ratios that was sampled by the concurrent schedules studied here, but the absolute reinforcement rates increase as the contours move away from the origin. Arranging concurrent schedules with the 37 combinations of scheduled absolute reinforcement rates depicted by the filled symbols in this panel is one way to sample more fully the domains of the single-alternative forms. Another way is by arranging concurrent schedules with the 54 combinations of scheduled absolute reinforcement rates indicated by the unfilled symbols in the top panel of Figure 5. To take one example from this panel, the symbol in the upper right corner represents a concurrent RI 1 RI 1 schedule. Obviously, the feedback properties of concurrent RI RI schedules will cause the obtained reinforcement rates on these schedules to shrink back toward the origin of the domain. The purpose of the experiment reported in the following section was to extend the evaluation of the evolutionary theory by more fully sampling the two-dimensional reinforcement rate domains of Equations 3, 4, 7, 8, and related forms.

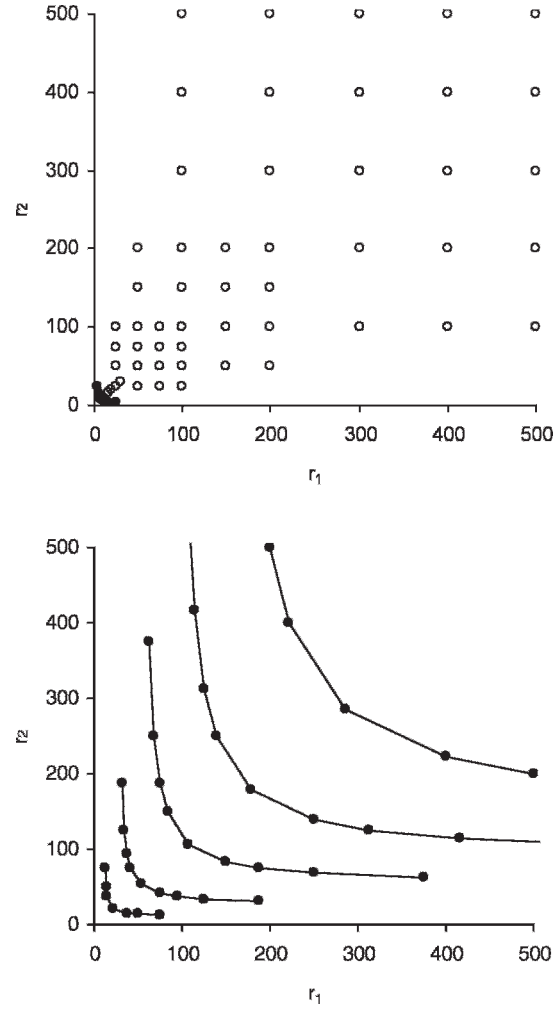

Fig. 5. Two-dimensional reinforcement rate domains for Equations 3, 4, 7, and 8, and related forms (e.g., Equations 7′ and 8′). Filled symbols in the top panel show the scheduled reinforcement rates in McDowell et al.'s (2008) concurrent schedule series. Each contour in the bottom panel samples the same range of reinforcement rate ratios that is specified by McDowell et al.'s concurrent schedule series, but the absolute reinforcement rates increase as the contours move away from the origin.

EXTENDED SAMPLING OF THE REINFORCEMENT RATE DOMAIN

In this experiment, the behavior of virtual organisms animated by the evolutionary dynamics was studied as they worked on 54 symmetric concurrent schedules at each combination of four psf mean pairs and five mutation rates. The scheduled reinforcement rates in the components of the 54 concurrent schedules are indicated by the open symbols in the top panel of Figure 5.

METHOD

Subjects, Apparatus, and Materials

The same subjects, apparatus, and materials used by McDowell, et al. (2008) were used in the present experiment. This included all details of the population of potential behav-

Table 2

RI components (in time ticks) of 54 concurrent schedules, corresponding to the scheduled reinforcement rates represented by the unfilled symbols in the top panel of Figure 5.

Schedule Number	Component		Schedule Number	Component	
	1	2		1	2
1	1.00	1.00	28	5.00	1.25
2	1.00	1.25	29	5.00	1.67
3	1.00	1.67	30	5.00	2.50
4	1.00	2.50	31	5.00	3.33
5	1.00	5.00	32	5.00	5.00
6	1.25	1.00	33	5.00	6.67
7	1.25	1.25	34	5.00	10.00
8	1.25	1.67	35	5.00	20.00
9	1.25	2.50	36	6.67	5.00
10	1.25	5.00	37	6.67	6.67
11	1.67	1.00	38	6.67	10.00
12	1.67	1.25	39	6.67	20.00
13	1.67	1.67	40	10.00	2.50
14	1.67	2.50	41	10.00	3.33
15	1.67	5.00	42	10.00	5.00
16	2.50	1.00	43	10.00	6.67
17	2.50	1.25	44	10.00	10.00
18	2.50	1.67	45	10.00	20.00
19	2.50	2.50	46	17.00	17.00
20	2.50	3.33	47	20.00	5.00
21	2.50	5.00	48	20.00	6.67
22	2.50	10.00	49	20.00	10.00
23	3.33	2.50	50	20.00	20.00
24	3.33	3.33	51	25.00	25.00
25	3.33	5.00	52	33.00	33.00
26	3.33	10.00	53	50.00	50.00
27	5.00	1.00	54	70.00	70.00

iors, such as circular wrapping of behavioral phenotypes, and division of the population into three classes, namely, two target classes consisting of 41 phenotypes each, and one extraneous class consisting of 942 phenotypes.

Procedure

Except for the specific concurrent schedules arranged, the procedure was identical to that used by McDowell et al. (2008) in every detail. This included the method of calculating fitness, the linear form of the parental selection function, the methods of bit string recombination and mutation, and the steps of the genetic algorithm that implemented the evolutionary theory, all of which are described in more detail in the Appendix. The 54 concurrent schedules arranged in the experiment are listed in Table 2. The entire set of schedules was arranged at each combination of four psf mean pairs, namely, 25-25, 40-40, 60-60, and 80-80, and five mutation rates, namely, 5, 7.5, 10, 12, and 20%, resulting in 20 experimental conditions. As in McDowell et al., at the start of exposure to each schedule, the population of potential behaviors was preloaded with roughly equal numbers of behaviors in the two target classes. This ensured that the probability of emitting the two target behaviors was relatively high and relatively equal at the start of exposure to a schedule. Each schedule remained in effect for 20,500 generations (time ticks) and each condition of the experiment was repeated 5 to 10 times. This produced over 100 million generations of responding.

Results and Discussion

As in previous research (McDowell, 2004; McDowell et al., 2008) reinforcements and responses in each component of a schedule were accumulated in 500-generation blocks. The first block was always discarded and frequencies in the remaining 40 blocks were averaged. The resulting response and reinforcement rates were then averaged across repetitions of a condition and these averages were used in all analyses.

As expected, the feedback properties of the schedules caused the obtained reinforcement rates to retract toward the origin of the two-dimensional reinforcement rate domain. The extent of this retraction was affected by the psf mean pair, and to a much lesser extent by the mutation rate. The larger the psf mean pair, the greater the retraction toward the origin. Specifically, for the 25-25, 40-40, 60-60, and 80-80 psf mean pairs, the obtained reinforcement rates did not exceed approximately 250, 200, 150, or 125 reinforcements per 500 ticks respectively, more or less regardless of mutation rate. As a result, the experimental conditions using the 25-25 psf mean pair effectively sampled roughly 25% of the reinforcement rate domain, the conditions using the 40-40 psf mean pair effectively sampled roughly 16% of the reinforcement rate domain, and the conditions using the 60-60 and 80-80 psf mean pairs sampled roughly 9% of the reinforcement rate domain. These sampling percentages greatly exceeded the roughly 0.25% sampling percentage for the data analyzed in the previous section.

Equations 7′, 8′, and 5′ of modern matching theory were fitted simultaneously as an ensemble to the 12 sources of variance at each of the five mutation rates. Hence, five ensemble fits of the modern theory's equations were carried out. The schematic diagram in Figure 2 again summarizes the analysis, which was the same as that used previously, except that each source of variance consisted of response and reinforcement rates from 54 rather than 11 concurrent schedules. For each ensemble fit, the sum of the 12 individual RSS/SS ratios (Equation 11) was minimized by adjusting one k parameter, shared across all sources of variance, and four sets of a, c_{1e} and c_{2e} parameters, each set shared by the three sources of variance at a specific psf mean pair. Five ensemble fits of the classic theory's Equations 3, 4, and 1 were also carried out by simultaneously adjusting one k parameter, shared across all sources of variance, and four r_e parameters, one shared by the sources of variance at each psf mean pair.

As in the previous analyses, %VAFs were calculated separately for each source of variance. These %VAFs ranged from 84% to 100% for the modern theory fits and from 71% to 96% for the classic theory fits. The median percentages of variance accounted for are plotted in Figure 6. The %VAFs for the classic theory did not include negative values and were much less variable than those shown for the classic theory in Figure 3, evidently because extreme mutation rates were not used in the present experiment. Overall, these %VAFs indicated that both the modern and the classic theory provided reasonably good descriptions of the data. Importantly, however, individual %VAFs for the modern theory exceeded those for the classic theory in all 60 comparisons of the two theories.

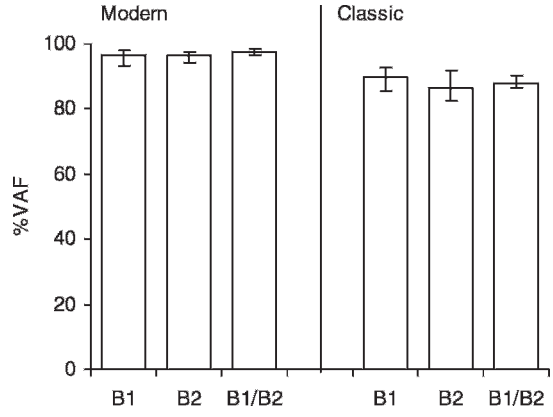

Fig. 6. Median percentages of variance accounted for (%VAF) by fits of Equations 7′, 8′, and 5′ of the modern theory, and Equations 3, 4, and 1 of the classic theory of matching to data from the extended domain sampling experiment. The x-axis labels identify the dependent variable of each equation. The error bars span the interquartile range.

The tradeoff between number of parameters and magnitude of the RSS for the two theories was assessed by means of the AIC_c, which was calculated as described earlier. The differences between the AIC_cs for the modern and classic theories were very large and favored the modern theory in all five comparisons. The large differences were again due to much larger RSSs for the classic theory fits. These differences indicate overwhelmingly that the modern theory described the data better than the classic theory.

The residuals left by the ensemble fits were pooled and plotted as before, and a cubic polynomial was fitted to each set of residuals. Significant ($\alpha = .05$) cubic polynomial trends were found for 4 of the 15 sets of the modern theory's residuals, and for all 15 sets of the classic theory's residuals. According to the binomial test used earlier, the probability of

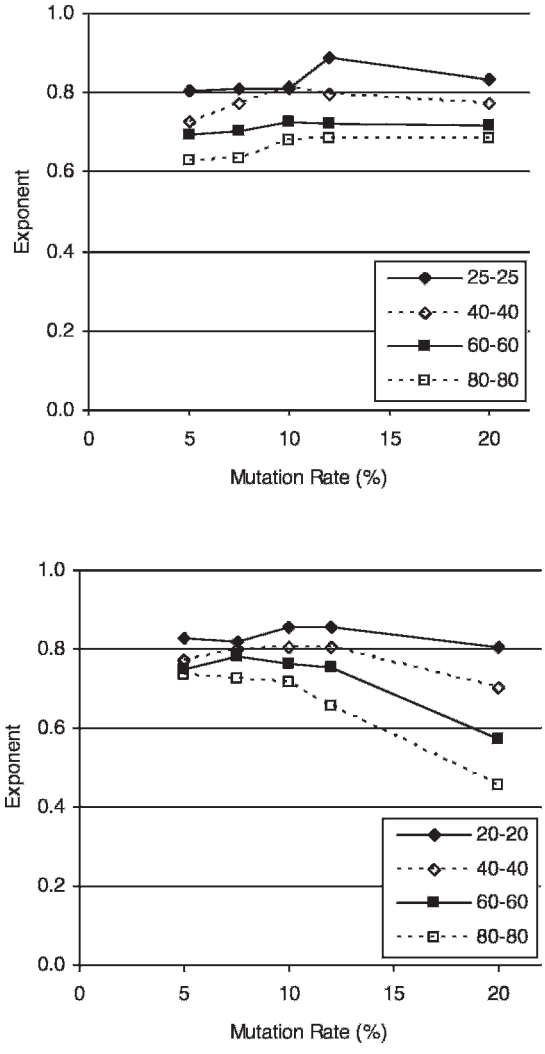

Fig. 7. Top panel: exponents from Table 3. Bottom panel: exponents from the corresponding conditions in the previous analysis (Table 1), where a more restricted reinforcement rate domain was used.

identifying 4 of 15 sets of residuals as nonrandom under the null hypothesis that all were random is .55, indicating that the residuals left by the modern theory can be considered random. For the classic theory, the probability of identifying 15 of 15 sets of residuals as nonrandom under the null hypothesis that all were random is essentially zero, indicating that the residuals left by the classic theory cannot be considered random. As in the previous analysis, the effect sizes of the cubic polynomial trends in the classic theory's residuals were often quite large, ranging from 0.14 to 0.90, with a median of 0.27. For the modern theory, the effect sizes of the cubic polynomial trends ranged from 0.00 to 0.08, with a median of 0.02. It might seem that an effect size within the latter range would not be judged statistically significant, but this occurred because of the large number of residuals in each set (216 = 54 concurrent schedules × 4 psf mean pairs), which translated into a large number of degrees of freedom for the error term in the trend test.

The parameters estimated from the ensemble fits are listed in Table 3. Exponents from the table are plotted in the top panel of Figure 7. All exponents ranged between 0.6 and 0.9 and tended to remain relatively constant across the mutation rates studied. This is in contrast to the exponents for the same conditions in the ensemble analyses carried out in the previous section. The latter exponents, listed in Table 1, are plotted in the bottom panel of Figure 7 for comparison. For the 25-25 (or 20-20) and the 40-40 psf mean pairs the exponents were roughly the same for the present data as for the data analyzed previously. But for the 60-60 and 80-80 psf mean pairs the exponents from fits to the present data tended to be a bit lower at the lower mutation rates than exponents from the previous analysis, and then remained at that level as mutation rate increased. In contrast, the exponents from the previous analysis tended to be higher at the lower mutation rates, and then fell at the higher mutation rates. Over the range of mutation rates studied in this experiment then, the behavior of the exponents differed depending on how extensively the reinforcement rate domain was sampled.

Bias favoring a concurrent schedule component over the background alternative was examined by calculating the quantity defined by Equation 12 and the analogous quantity for b_{2e}/r_e. As was the case for the previous analysis (illustrated in the bottom panel of Figure 4), these quantities declined as the selection events in the component became weaker relative to the background alternative. However, the absolute magnitudes of the quantities containing the bias parameters were always larger for the present data than for the data analyzed previously. Finally, as shown by the c_{2e}/c_{1e} quantities in Table 3, there was no bias favoring one component over the other in these symmetrical concurrent schedules.

Table 3

Parameters of the equations of modern and classic matching theory estimated from ensemble fits to data from the extended domain sampling experiment.

		Modern					Classic	
psf mean pair		k	a	c_{1e}	c_{2e}	c_{2e}/c_{1e}	k	r_e
	5%							
25-25		284.46	0.80	1.62	1.58	0.98	244.19	0.12
40-40			0.73	5.58	5.38	0.96		12.44
60-60			0.69	12.67	12.73	1.00		26.42
80-80			0.63	14.90	15.90	1.07		44.68
	7.5%							
25-25		285.67	0.81	3.38	3.41	1.01	244.12	1.78
40-40			0.77	9.08	9.17	1.01		19.70
60-60			0.70	15.67	15.85	1.01		32.20
80-80			0.63	17.73	17.61	0.99		52.96
	10%							
25-25		289.21	0.81	5.66	5.59	0.99	250.70	4.83
40-40			0.82	14.19	14.21	1.00		29.08
60-60			0.73	20.28	20.24	1.00		42.84
80-80			0.68	24.39	24.01	0.98		63.99
	12%							
25-25		285.39	0.89	9.16	9.11	1.00	245.06	5.71
40-40			0.80	14.68	14.92	1.02		31.48
60-60			0.72	21.29	21.32	1.00		44.92
80-80			0.69	26.05	25.92	1.00		66.02
	20%							
25-25		300.11	0.83	17.75	17.51	0.99	240.72	14.87
40-40			0.77	25.44	25.22	0.99		48.31
60-60			0.72	32.44	31.99	0.99		64.12
80-80			0.68	37.76	37.88	1.00		85.33

Note. The parameter, c_{2e}/c_{1e}, was calculated from the estimates of c_{1e} and c_{2e}. The spanner heads identify the mutation rate.

These results showed that when the absolute reinforcement rate domain was sampled more fully, the behavior generated by the evolutionary dynamics remained consistent with the modern theory of matching, including its constant-k assumption, and remained inconsistent with the classic theory of matching. However, the results also showed that some properties of the behavior generated by the evolutionary theory differed when the reinforcement rate domain was sampled more extensively. This difference was reflected in differences in the exponent and the bias parameters of the modern matching theory equations.

GENERAL DISCUSSION

The evolutionary theory of behavior dynamics consists of low-level rules of selection, reproduction, and mutation that are formally unrelated to higher-level, time-averaged properties of the behavior the rules generate. Consequently, it is remarkable that the higher-level behavior agreed precisely with the modern theory of matching. This agreement was observed in ensemble fits that enforced the theory's constant-k requirement, and also enforced at each psf mean pair the identity of parameters theoretically (i.e., mathematically) required to be identical across the three equations of the theory. These restrictive fits not only accounted for large percentages of response rate variance, but they also left random residuals, which indicates that no other account is likely to describe behavior generated by the evolutionary theory better than modern matching theory.

Of course, the agreement between the evolutionary dynamics and modern matching theory constitutes support for the former only if the latter is the correct account of the steady-state behavior of live organisms. As noted earlier, McDowell (2005) argued that this

account was tenable, while attempting to show that the classic theory was false. Although few data sets have been subjected to the restrictive ensemble fits conducted here, whenever such fits have been reported (Dallery et al., 2004; McDowell, 2005), the results were consistent with the modern theory of matching, and inconsistent with the classic theory. Nevertheless, further testing is surely worthwhile, especially with regard to the constant-k assumption of the modern theory. The analogous assumption of the classic theory proved to be its downfall, according to some researchers (Dallery & Soto, 2004; McDowell 2005), and consequently its modern theory version merits careful scrutiny.

The values and properties of the parameters from the ensemble fits of the modern matching theory equations lend further support for the evolutionary dynamics. At moderate mutation rates, exponents, a, ranged between about 0.6 and about 0.9, a range that is often observed in experiments with live organisms. But note also that when the equations were fitted to individual repetitions of conditions at moderate mutation rates in McDowell et al. (2008), the exponents varied more widely, including values as small as about 0.3 and values larger than unity, with a mean and median remaining at about 0.8. It is not uncommon to observe these more extreme values in experiments with live organisms.

The bias parameters, b_{1e}, b_{2e}, and b_{12} ($= c_{2e}/c_{1e}$) estimated from these fits also behaved in ways that were consistent with bias parameters obtained from experiments with live organisms. Specifically, the first two parameters decreased as the selection events in the concurrent schedule components weakened relative to the unchanging background alternative, and b_{12} varied around unity in the symmetric concurrent schedules.

Regarding the evolutionary theory itself, it is essential to obtain an accurate understanding of its emergent outcomes, both for the sake of comparing them to data from live organisms, and also for the sake of making reliable novel predictions that can be tested in live experiments. It is important that the outcomes and predictions in fact follow from the theory and are not artifacts due to, for example, inadequate sampling of independent variable domains. The results shown in Figure 7 are a case in point. One may ask what the true behavior of these exponents is, according to the evolutionary theory. When the entire reinforcement rate domain is sampled, would the exponents remain constant at more extreme mutation rates, as is perhaps suggested by the results plotted in the top panel of Figure 7, or would they decline as suggested by the exponents in the top panel of Figure 4? A third possibility is that they would remain relatively constant when the entire domain is sampled, but would decline markedly at mutation rate extremes whenever one of the ''restricted'' contours shown in the bottom panel of Figure 5 is followed. These are of course questions that can only be answered by further computational experiments.

A puzzling feature of the evolutionary theory is what, if anything, its elements might be understood to represent in the material world. Clearly, behaviors do not have genotypes, cannot reasonably be said to mate, or to reproduce by recombination or in any other way, and so on. Of course the theory can be understood on purely formal grounds, just as an equation can be understood as describing events in the natural world without itself corresponding to anything material. However, it has been argued elsewhere (McDowell, 2004) that the accuracy of the evolutionary theory suggests that whatever material events are responsible for instrumental behavior, they must be computationally equivalent to selection by consequences. McDowell and Caron (2007) noted that this view, when taken to its logical conclusion, suggests that organic evolution produced a copy of itself in the nervous systems of biological organisms for the purpose of regulating their behavior during their individual lifetimes.

Interestingly, a selectionist theory of brain function was proposed more than 20 years ago by Gerald Edelman (1987). This is a theory of whole brain functioning in intact organisms behaving in environments where actions can have consequences. Briefly, the theory proposes that groups of neurons subserving behavior are selected by these consequences, and thereafter are more likely to become activated under similar circumstances. Edelman's theory, which needless to say is more extensive and complex than suggested here, is modeled after his successful selectionist theory of immune system function, for which he won the Nobel Prize in Physiology or Medicine in 1972.

Edelman (2007) recently implemented a version of his theory in a mechanical agent, known as Darwin VII, which showed rudimentary adaptive behavior in a laboratory environment. McDowell (2010) examined more fully the relationship between the evolutionary theory of behavior dynamics and Edelman's theory, and noted that

> …the bit, the bit sequence, the bit string, and the bit string class are realized in [Edelman's theory] as a neuron, a synapse, a neuronal group, and a collection of degenerate neuronal groups [all having the same function]. In addition, the action of selection, which is carried out formally by the parental selection function in the [evolutionary theory of behavior dynamics] is realized in [Edelman's theory] by the operation of diffuse value systems in the brain that alter synaptic strengths or thresholds.

The evolutionary theory is a functional dynamic theory in the same sense that, to take a well known example, Newton's second law, $F = ma$, is a functional dynamic theory. Newton's second law describes in a completely general way how an unbalanced force causes an acceleration. To understand how the general law applies to a specific material object, say, a rocket nozzle, requires detailed knowledge of the nozzle's properties, the physics of expanding gasses, and so on. Similarly, the general evolutionary theory of behavior dynamics may be implemented in biological organisms by means of something like the selectionist neural mechanisms proposed by Edelman. If so, the result would be a comprehensive account of adaptive behavior that extended from brain function to the behavior of whole organisms in quantitative detail.

REFERENCES

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22*, 231–242.

Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior, 32*, 269–281.

Baum, W. M., & Rachlin, H. C. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior, 12*, 861–874.

Dallery, J., McDowell, J. J., & Soto, P. L. (2004). The measurement and functional properties of reinforcer value in single-alternative responding: A test of linear system theory. *Psychological Record, 54*(1), 45–65.

Dallery, J., & Soto, P. L. (2004). Herrnstein's hyperbola and behavioral pharmacology: Review and critique. *Behavioural Pharmacology, 15*, 443–459.

Dallery, J., Soto, P. L., & McDowell, J. J. (2005). A test of the formal and modern theories of matching. *Journal of the Experimental Analysis of Behavior, 84*, 129–145.

Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection.* New York: Basic Books.

Edelman, G. M. (2007). Learning in and from brain-based devices. *Science, 318*, 1103–1105.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267–272.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*, 243–266.

McDowell, J. J (1986). On the falsifiability of matching theory. *Journal of the Experimental Analysis of Behavior, 45*, 63–73.

McDowell, J. J (2004). A computational model of selection by consequences. *Journal of the Experimental Analysis of Behavior, 81*, 297–317.

McDowell, J. J (2005). On the classic and modern theories of matching. *Journal of the Experimental Analysis of Behavior, 84*, 111–127.

McDowell, J. J (2010). Behavioral and neural Darwinism: Selectionist function and mechanism in adaptive behavior dynamics. *Behavioural Processes, 84*, 358–365.

McDowell, J. J, & Caron, M. L. (2007). Undermatching is an emergent property of selection by consequences. *Behavioural Processes, 75*, 97–106.

McDowell, J. J, Caron, M. L., Kulubekova, S., & Berg, J. P. (2008). A computational theory of selection by consequences applied to concurrent schedules. *Journal of the Experimental Analysis of Behavior, 90*, 387–403.

McDowell, J. J, & Dallery, J. (1999). Falsification of matching theory: Changes in the asymptote of Herrnstein's hyperbola as a function of water deprivation. *Journal of the Experimental Analysis of Behavior, 72*, 251–268.

McDowell, J. J, & Popa, A. (2009). Beyond continuous mathematics and traditional scientific analysis: Understanding and mining Wolfram's *A New Kind of Science. Behavioural Processes, 81*, 343–352.

Motulsky, H. J., & Christopoulos, A. (2004). Fitting models to biological data using linear and nonlinear regression. New York: Oxford University Press.

Soto, P. L., McDowell, J. J., & Dallery, J. (2005). Effects of adding a second reinforcement alternative: Implications for Herrnstein's interpretation of r_e. *Journal of the Experimental Analysis of Behavior, 84*, 185–225.

Staddon, J. E. R. (1968). Spaced responding and choice: a preliminary analysis. *Journal of the Experimental Analysis of Behavior, 11*, 669–682.

Received: September 1, 2009
Final Acceptance: February 3, 2010

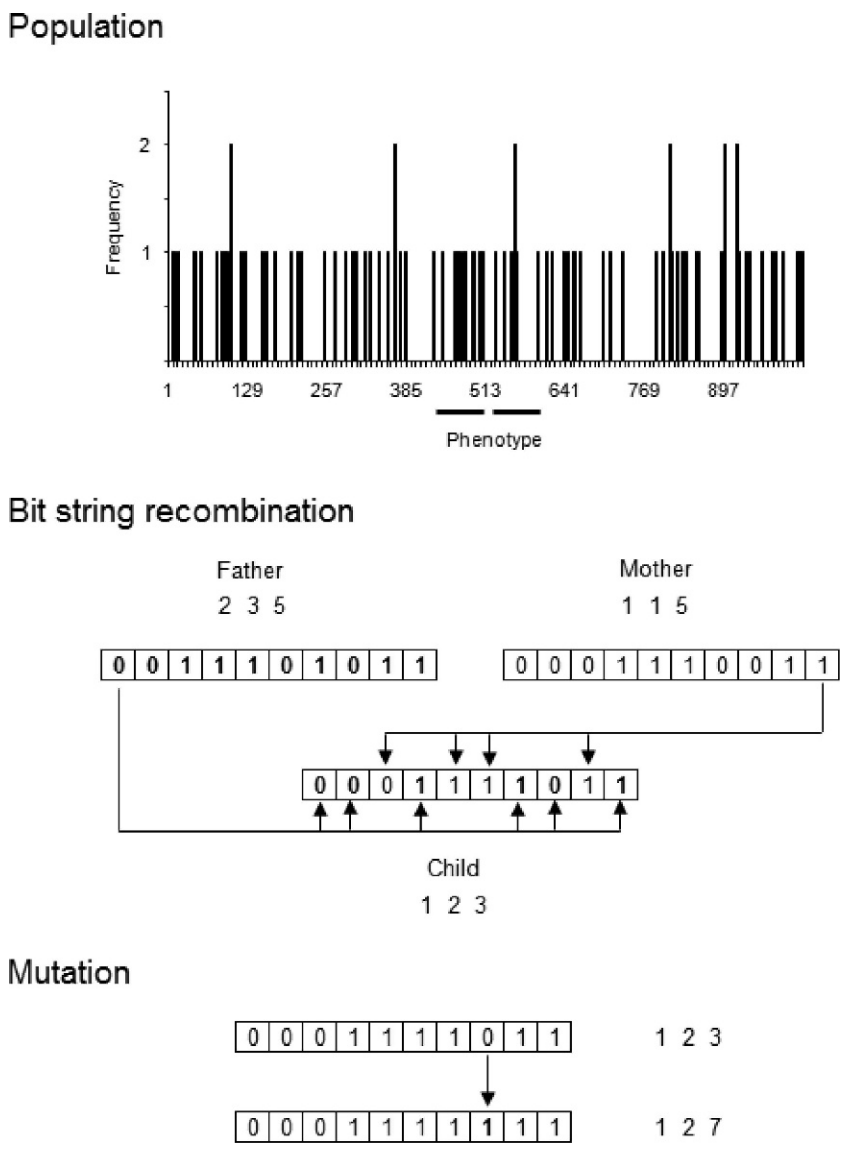

Fig. A1. Top panel: a frequency distribution of 100 potential behaviors selected at random from the range, 0 to 1023. The integer values are the behaviors' phenotypes. The horizontal bars beneath the *x*-axis span two operant classes. Middle panel: father and mother bit strings are recombined to produce a child bit string. Bits in a child's bit string come from the father (bold) or the mother with equal probability. Decimal integers (phenotypes) corresponding to each bit string (genotype) are given in the panel. Bottom panel: mutation occurs in a fraction of a new population's behaviors when a random bit in the behavior's bit string is flipped.

Appendix

The selectionist theory operates on a population of 100 potential behaviors, each of which is represented by a 10-character sequence of 0s and 1s. These ten-character bit strings range from ''0000000000'' to ''1111111111'' and constitute binary representations of the decimal integers, 0 through 1023. A behavior's bit string is its genotype; the decimal integer into which the bit string decodes is its phenotype. The top panel of Figure A1 shows a frequency distribution of 100 behavioral phenotypes drawn at random from the range 0 through 1023. This is a population of potential behaviors.

Behavioral phenotypes can be grouped into operant classes. For all concurrent-schedule experiments discussed in this article, two target classes were defined, one by the 41 integers from 471 through 511, the other by the 41 integers from 512 through 552. These two ranges are indicated, approximately, by the horizontal bars beneath the *x*-axis in the top panel of Figure A1. The two ranges of phenotypes correspond to behaviors that successfully operate two manipulanda. Behaviors that do not fall in either operant class are extraneous behaviors.

At each moment of time a behavior is randomly emitted from the population of potential behaviors that exists at that moment. Following emission, 100 pairs of parent behaviors are chosen from the population and are mated to produce 100 child behaviors, which constitute the next generation. If the emitted behavior was reinforced, then the parents are chosen on the basis of their fitness. If the emitted behavior was not reinforced, then the parents are chosen at random. In all experiments discussed in this article, the fitness, *x*, of a behavior was defined as the absolute value of the difference between its phenotype and the middle phenotype of the target class from which the reinforced behavior was emitted. Hence, smaller values of *x* corresponded to fitter behaviors. The fitness of the behaviors in the population was then used to choose parents according to the linear probability density function,

$$p(x) = -\frac{2}{9\mu^2}x + \frac{2}{3\mu}, \qquad \text{(A1)}$$

where μ is the mean of the density function, and $0 \leq x \leq 3\mu$, which means that behaviors with fitness values $\geq 3\mu$ had no chance of becoming parents. This is the simplest linear density function that depends only on its mean. Note the negative slope. Every behavior in the population is assigned a fitness, *x*. The greater its fitness (i.e., the smaller the *x*), the greater the probability density associated with its becoming a parent. Parents chosen in this way are relatively fit, that is, relatively close in integer value to the emitted behavior. Notice also that the average fitness of parents chosen in this way is inversely related to the mean of the density function. In the text, linear density functions with various means are referred to as parental selection functions.

After 100 pairs of parents are chosen, each pair is mated to produce one child behavior. Mating is illustrated in the center panel of Figure A1, which shows a form of multipoint crossover recombination of the parents' bit strings. In all experiments discussed in this article, each bit in the child's bit string had a 50-50 chance of coming from the same location in the father's or the mother's bit string. This method of recombination produces child behaviors with phenotypes that are more or less similar to the father's, the mother's, or both parents' phenotypes.

The new population produced by recombination of the parent's bit strings then undergoes a small amount of mutation, which is illustrated in the bottom panel of Figure A1. In all the experiments discussed in this article, each behavior in the new population had a finite probability of undergoing mutation (referred to in the text as the mutation rate), which consisted of flipping one randomly chosen bit.

Following mutation, a randomly chosen behavior from the population is emitted and the processes of selection (if the behavior is reinforced), reproduction, and mutation are repeated. This generates a continuous stream of behavior that can be studied just as if it were produced by a live organism. The elements described in this appendix constitute the selectionist theory in its entirety. It has no other features, properties, or constraints.